Before the

Federal Communications Commission

Washington, D.C. 20554

In the Matter Of	)	
A National Broadband Plan for Our Future	)	GN Docket 09-51

COMMENTS OF Ionary Consulting

The Commission now has an opportunity of historic proportions to reshape the nation's communications infrastructure. Since 1934, its narrow focus has largely been to create "universal service" for the analog wireline network. But the world has moved beyond that stage. Between higher-speed services and especially the Internet, and the widespread use of wireless services, the old priorities must change. Regulatory structures must be reshaped to reflect the new reality. We are concerned that the current framing of "broadband" is not workable, and that the questions in the Commission's Inquiry must first be set into a proper context. Hence our Comments will focus on selected issues, after first discussing how the issues should be framed.

Ionary Consulting is a proprietorship based in Newton, Massachusetts that provides consulting services to competitive service providers (CLECs, ISPs, cable and wireless) and their suppliers. Its principal, Fred Goldstein, has over 30 years of experience working with the telecommunications industry.

First, properly frame the issue

The issue under discussion – A National Broadband Plan for our Future – must be framed correctly. The Commission over the past eight years essentially gutted over two decades of pro-competitive, pro-consumer policies. It destroyed workable elements of the regulatory structure and created instead an unworkable mess. This Inquiry includes hundreds of questions. While they could be answered individually, a more productive approach is to create a proper framing for the topic in general, one in which the answers to many specific questions fall out naturally.

It is thus our position that the first stage of a "broadband" plan is to create a frame of reference, to define what we mean by "broadband". We posit that the word is *not* a simple noun. Trying to treat it as a single thing can only result in failure; the way that the thing is defined necessarily dictates much of how it is regulated. We instead posit two alternative ways to view the term, both of which support the same framing of the issues.

One approach is to treat "broadband" as an adjective. This is the grammatically-correct approach. Like a trademark (in many opinions, properly an adjective), it should be used to describe some thing. In this sense, "broadband" as commonly used today describes two very different things.

One thing that can be described by the adjective "broadband" is a transmission medium. "Broadband" began as a term for frequency-division multiplexed signals over coaxial cable, as used in the CATV industry; hence the little-used 1980s Ethernet standard $10_{broad30}$. The CCITT used it to refer to bit rates over 50 Mbps, hence "Broadband ISDN", the project that begat Asynchronous Transfer Mode. Nowadays, "broadband" is used to describe a transmission medium *capable of* carrying a high-bit-rate signal of some sort. So a digital subscriber line or a cable modem are broadband media, providing *broadband transmission*.

The second usage of the adjective is to refer to specific content *services* that are carried atop broadband transmission. Broadband services may include Internet access, of course, but could also include other services such as switched digital video and private data network access.

An alternative approach is to draw linguistic analogy to the two Spanish nouns "el radio" and "la radio". The first refers to the device (hardware); the second refers to programming (software). There is thus a distinction between "el broadband" (a high-speed transmission medium) and "la broadband" (services and content carried across a high-speed medium). What we have had since 2005, on the other hand, is a vertically integrated regulatory model, one in which the content is essentially the property of the transmission-medium owner. This policy has failed, leading instead to increased conflict.

In the 1930s, when the Communications Act was passed, the monopoly telephone network was predominantly used for one purpose, voice calls, This one Service drove the design of the network and the entire plant. AT&T was vertically integrated from wire mill to phone bill. But the 1968 Carterfone Decision, 1968-1988 Computer Inquiries, 1975-1982 MTS/WATS

proceedings and the Telecom Act of 1996 all led to a more open industry, with clearer distinctions between the roles of each player. All of that progress is lost when "broadband" is treated as one unified thing.

The first stage of any Plan, then, must restore the distinction between broadband transmission and broadband service. It is the distinction between carriage and content, between carrier and shipper, between highway and vehicle. The first question that then arises is how carriage (transmission as a service) should be provided. If carriage is restored to the kind of wide-open bit-neutral medium that it was just a few years ago, then market forces and competition will necessarily take care of the rest. Most regulation of service and content providers becomes redundant.

Specific Broadband Service Questions

We now address selected specific questions raised in the Notice of Inquiry.

In its Notice (at 16), the Commission asks,

> For instance, the Commission currently uses the terms "advanced telecommunications capability," "broadband," and "high-speed Internet." Should these definitions be unified, or should they have separate meanings for different purposes, keeping in mind that current and future broadband platforms will increasingly support "high-speed Internet" as one of several offered services including voice, video, private data applications, and the like?

It is clear from our framing of the question that these terms describe very different entities. An "advanced telecommunications facility" is a lower-layer transmission (carriage) entity. "Internet" refers to a type of payload carried atop carriage of some sort. As noted in the Attachment, Internet is actually a business model for networks, not a single entity, though there is one "big-I Internet" that currently demonstrates the potential of the concept. Video and voice are most often carried separately from the Internet, although it is of course possible to carry some video and voice across the Internet. The PSTN is not the Internet, and a Title V Cable Service is not the Internet, but these can all be carried across broadband transmission facilities. Treating all of them as one can only cause problems.

At 17, the Commission asks "whether a definition of 'broadband' should be tethered to a numerical definition or, instead, an 'experiential' metric based on the consumer's ability to access sufficiently robust data for certain identifiable broadband services." At 18, a complementary question is raised: "We also request comment on whether a definition of broadband should be static or dynamic, with speed tiers that adjust with changes in technology." We question the value of attempting to place a bright-line definition over the term "broadband". Clearly any attempt to define it should be dynamic, as both technology and consumer expectations change over time and, quite possibly, from place to place.

But the key is not to define "broadband" carriage as a special case, different from other facilities and carriage. In recent years, ILECs in particular have used "broadband" as an excuse to evade their Telecommunications Act obligations, as if somehow Section 251 were limited to voice-grade services and Section 706 was meant to exempt "advanced" services from every other section. Such interpretations are fanciful at best and have proven malignant in practice. Incumbent carriers should be required to offer broadband common carriage and to offer unbundled facilities suitable for the competitive provision of broadband services, while non-common-carrier service providers (such as wireless ISPs and cable providers) should be permitted to continue to self-provision their services, such as Internet access, regardless of arbitrary definitions.

Many questions are answered by simply accepting that the Telecommunications Act's pro-competitive changes apply at all speeds "from DC to daylight". It is frankly nonsensical to waive obligations on transmission services at 24 Mbps that apply at 2.4 kbps. "Carrier of last resort" obligations, as well as associated subsidy programs, should apply equally at all speeds. Waivers of regulatory obligations were issued in the early 2000s as an incentive for carriers to construct facilities, as if they would not build them if subject to the clear black-letter requirements of the Telecommunications Act. Most other countries – virtually the entire OECD – did not take that approach, and retained some unbundling and common carriage obligations for advanced ("broadband") facilities. During that time, the United States fell in its international position Deregulation backfired.

A reasonable view must be taken of what can be competitive

At 25, the Commission asks,

> To what extent should the Commission consider price or marketplace competition for broadband as it considers whether people have access to broadband capability? For example, how should the Commission consider the benefits of consumers in a particular area having only a single provider, using one type of technology, versus the competitive benefits that could result from having one or more providers using similar or different technologies?

This is precisely the type of question whose answer is dictated by the framing of the question. Broadband transmission ("el broadband") is naturally much less competitive than broadband content and services ("la broadband"). In many locations – indeed, this may be the general case – broadband transmission is a natural monopoly. As such, the incremental cost to the incumbent of providing additional service is lower than the cost borne by a new entrant. This makes competition very difficult. The cable/ ILEC duopoly, rather than a pure monopoly, exists for two main reasons. One is that the two industries had technologies that did not overlap in capability until the 1990s, and thus each required its own specialized plant. The other is that the FCC had the foresight, in the 1970s, to prohibit telephone companies from owning the cable company in all but the highest-cost rural areas.

The level of competition at the transmission layer is thus not critical *if* the transmission facilities are themselves adequate *and* they are available on a common carrier basis, at just and reasonable rates, to all comers, or can be leased on an unbundled basis by any qualified CLEC. In either case, upgrades to the natural monopoly outside plant to improve its broadband capability is made *more* feasible by compelling these two wholesale obligations. This is the best meaning of "open" networks, transmission facilities that are available on a wholesale basis. When this is done, there is no need to regulate most content or services; they are not a natural monopoly, and market forces can determine such touchy questions as "what constitutes reasonable network management?"

It is when there are only one or two providers of broadband *services* that these service providers have any realistic chance of interfering with the market-dictated degree of "neutrality" that any given service should provide. If there is open broadband transmission, then that natural-monopoly plant can be used by new entrants at any time. If all of the existing service providers adopt policies that their customers do not like, then new entrants will naturally gravitate towards

that market *if* the customers' desires can be met economically. If no provider steps forward, then it may well be that the customers' desires simply cannot be met by current technology at reasonable prices. This is a perfectly rational market dynamic.

In fact, this is precisely how the Internet service market did work *until 2005*, when the FCC withdrew the common carrier obligation from ILECs, compounded by earlier reductions in unbundling obligations that would have allowed CLECs to fill in the gaps. While the facility-owners' captive ISPs had already taken the lions' share of retail customers, there were still credible alternatives, and they provided a check on the facility-owners' behavior. The sudden removal of the ILECs' common carrier and carrier of last resort status completely changed the market dynamic. Rather than have a large number of possible service providers, most areas had only two, some areas only one, and no open entry to new providers. The market dynamic of a two-provider market is very different from one with open entry.

It is thus incumbent upon the Commission to ensure that *at least one* broadband transmission provider *with common carrier obligations* is available wherever possible, so that there can be *many* broadband service providers. This is the direction taken in most countries. We note the recent Australian initiative, to have a nationwide network built if necessary by public funds, providing wholesale access on an open basis to retail service providers. We also note that the transmission plant in the United Kingdom has been functionally separated from the retail ILEC. These are examples of ways to structurally ensure that open access to service providers remains intact even over a monopoly or duopoly physical plant.

Wireless service is not a natural monopoly; the number of providers is set by the FCC in its licensing policy. In some areas, especially rural, wireless may be the only economically-feasible way of delivering broadband services. In urban areas where wireline broadband facilities are available, wireless is unlikely to be competitive except for customers with significant nomadic or mobile requirements. This is simply because the total capacity of the radio spectrum currently assigned to these functions is too low to meet demands, so it is relegated to specialized applications that cannot be met by wireline.

In low-population-density rural areas, wireless spectrum may well suffice to meet demand for broadband services, *provided* that it can be efficiently made available. This is not always the case today. Much or most of the currently-licensed spectrum, especially in the 2.3 GHz (WCS)

and 2.6 GHz (BRS) bands, is not being utilized. It is often being "banked" by major CMRS carriers in order to prevent competitors from using it. In some cases, even in the more established PCS band, a license area includes a major city and large rural areas, and the licensee focuses on serving the more-profitable urban areas. The Commission should consider fine-grained "use it or lose it" policies to encourage the utilization of spectrum in rural areas, and to make it available to local providers when licensees do not serve all of their licensed area.

What technologies are most appropriate in rural and rustic areas?

The Commission asks, at 38,

> Should the national broadband plan seek to bring broadband to 100 percent of the country? If so, what are the costs and benefits of bringing broadband to the least densely populated areas? We seek comment on how we can better estimate the cost of deploying various alternative broadband technologies to those areas that the market is not serving, or not adequately serving. Which broadband technologies might work best and deliver the most effective, efficient services in various parts of the nation?

This is a very important question. The best technology for an urban area is unlikely to be best in rural areas. Population density impacts cost. It is obvious that urban areas can be rewired with fiber optic more cheaply than rural ones, measured on a per-home basis. Much of the cost is per mile, and densely-populated areas have more homes per mile, hence a larger divisor.

In order to quantify the problem, we used a set of 5-digit ZIP code tabulation area population estimates (for 2005) to divide the 49-state land mass (omitting Alaska) into six density categories. This is more granular than common county-based maps and thus provides somewhat more insight into the population distribution.

Density Cell by ZIP/ pop. sq. mi	ZIPs	SqMi	kpop	Avg. pop/sq.mi	% of area	% of pop
cell 4 - dense urban (>7000)	937	2,744	34,647	12,629	0.1%	11.7%
cell 3 – urban (700-7000)	5,341	65,933	131,081	1,988	1.9%	44.3%
cell 2 – suburban (70-700)	7,412	470,336	88,475	188	13.4%	29.9%
cell 1 – rural (7-70)	11,857	1,790,285	38,878	22	50.9%	13.1%
cell 0 - rustic excl. Alaska (<7)	4,323	1,190,610	2,646	2.2	33.8%	0.9%
total	29,870	3,519,907	295,727	84	111.5%	
cell 9 - isolates (subset of rustic)	262	406,512	1,458	3.6	11.5%	

This is more granular than the common urban/suburban/rural division. The dense urban category applies to a number of urban cores with collectively little land; many of these have a high business density and thus teledensity disproportionate to their population. Along with urban and suburban areas, about 85% of the population lives in areas where fiber optics are likely to be economical if provided on a shared, wholesale basis (that is, a single provider can expect to have the preponderance of the market). In the lower-density suburban areas, DSL may be more appropriate in the interim, and may have a longer lifespan where the per-subscriber cost of new facilities is high and the existing copper plant can be upgraded at lower cost. To a great extent this varies on a case-by-case basis.

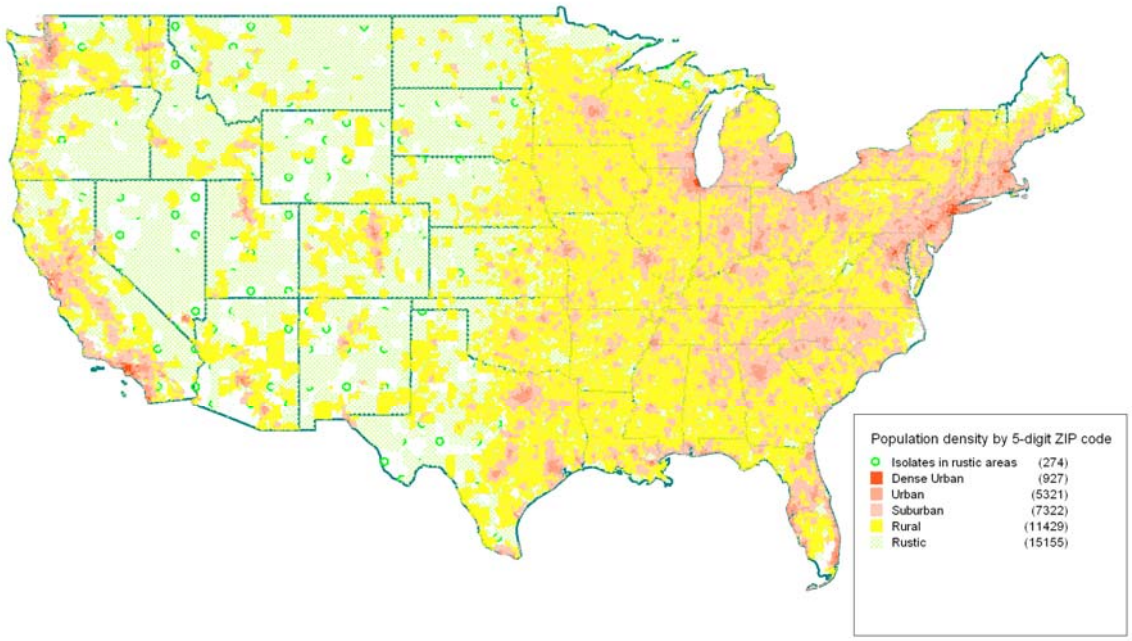

At the other extreme, the rustic category applies to areas that are unpopulated or have a lower population density than the typical rural farmland of the midwestern and southern states. Almost half of the total land area is rural, but the approximately one-third of the land area that is rustic has less than one percent of the total population. This includes mountain and desert areas, as well as Plains areas where agriculture is accomplished with relatively few inhabitants. The isolates category is a subset of rustic that includes ZIP code tabulation areas whose overall density is low due to a large area, but which have over 2500 in total population. Most of these isolated

concentrations are scattered in small towns and villages; they account for over half of the rustic population. As such, most of their population may be easier to serve than the scattered population of other rustic and even many rural areas.

One conclusion is that there is a small fraction of the population, in rustic areas outside of the isolated concentrations, that would be disproportionately expensive to reach using any available broadband technology, though some could be reached via terrestrial radio and most can use satellites (which are admittedly not a good substitute). While these locations should not be entirely written off, the cost of serving them is disproportionately high, and subsidies to them could disproportionately tap out financial resources. The best bet for those areas might be to make unlicensed spectrum available, with higher power limits than in other areas, and encourage innovative development of techniques such as wide-coverage mesh networks.

Many rural areas are characterized by small towns whose population densities are actually higher than many suburban markets. Wireline (coax, copper and fiber optic) facilities are thus useful for the last mile in these areas, just as in larger cities, but the backhaul (middle mile) cost is more likely to be the gating issue.

In most of the rural areas outside of these population centers, the density is still high enough to support a radio-based system. If a radio system (e.g., a cell) supports a 10-mile radius, then with the average rural (as defined above) population of 22 per square mile, a cell still averages 6800 people in its coverage radius. Of course real-world rural coverage may be more difficult due to terrain and other issues, but in general the range of radio systems is adequate. 4G radio technology such as LTE and WiMAX has superior capacity to older systems, and with their intelligent antennas, can provide both superior range and higher speed (to close-in users) than today's systems.

But this does not mean that current CMRS systems, *as they are typically operated*, are an adequate form of "broadband". Virtually all current CMRS data offerings are so tightly restricted that they bear little semblance to the familiar Internet services offered over DSL and cable. These "wireless web" and Deep Packet Inspection-laden services are simply not adequate. A properly *open* radio system should provide a broadband transmission medium that could be leased on a wholesale basis by broadband service providers, such as ISPs. Because wireless *bits* are more expensive than Wireline bits, on an incremental basis, the pricing structure of these

services might differ; today's common "all you can eat" Internet services may not be affordable, and block-of-byte rates might be the more usual retail offering. Offering unlimited access to a small range of applications (e.g., common "web and mail only" plans) is not adequate. Perhaps licensing policies could be adjusted to encourage a more open usage policy.

Small wireless ISPs, generally operating on an unlicensed basis, a non-exclusive basis on the 3650 MHz band, or with small spectrum leases in the 2300 and 2600 MHz bands, should not be subject to similar rules. They rarely operate restrictive services similar to those of the CMRS providers. WISPs have never been common carriers, do not (except for a few) have exclusive rights to spectrum, and tend to operate on a low budget. Their continued growth should be encouraged as a complementary strategy that provides some useful competitive pressure on other providers. A wholesale requirement might, however, be appropriate for those who take subsidy money and offer to take "broadband carrier of last resort" status upon themselves.

High-cost support should be for facilities, not services

At 39, the Commission asks,

> We seek comment on the impact of broadband on our existing universal service programs, and how we should conduct our analysis of the High-Cost, Schools and Libraries, Rural Health Care (including the Rural Health Care Pilot program), and Low-Income programs. Specifically, for each program, we seek comment on the program's effectiveness and efficiency as a mechanism to help achieve national broadband goals.48 Further, we seek comment on what modifications to these programs, if any, should be considered as a part of a national broadband plan. We seek comment on how these programs might be better targeted to address broadband deployment, particularly because these programs treat the support of broadband differently.

And at 41 the related issues continue,

> Should we modify existing universal service programs? For example, should we make broadband a "supported service" eligible to receive support directly from the High-Cost and Low-Income programs?...

Rather than comment on each program individually, we focus our comment on the way high-cost support is allocated. This program is fundamentally backwards-looking, as it focuses on

minimizing the retail price of local telephone service in rural areas. It thus subsidizes service providers, who are free to construct whatever facilities they wish. This has proven to be essentially a blank check for rural ILECs, some of whom have built gold-plated FTTx networks. The equal cost rule has then allowed some CETCs to operate in their territories and collect equivalent per-line subsidies, regardless of cost. The Commission's 2008 proposal was to phase out CETC subsidies and simply cap ILEC ETC subsidies. This would limit the growth of USF, to be sure, but it does not solve the problem.

In rural and rustic areas, wireline plant is even more clearly a natural monopoly than it is in urban areas. The Commission has allowed rural ILECs to provide the cable service in their areas, since separate cable operations would be unprofitable. And yet many rural ILECs have been allowed to drop their common carrier status, creating a broadband *service* monopoly.

Adding broadband as a "supported service" leaves the same structure in place. When telephony was the only major service, it could be directly supported. But "broadband" is not one service. It describes a range of services, as described above. Thus the focus should not be on services per se.

The more appropriate basis of USF is to subsidize the *facilities* that carry broadband, such as the local loops, fiber-to-the-premise, and middle-mile backhaul facilities. Subsidized carriers should then be required to provide these facilities on an unbundled basis, while retaining their USF subsidies on those lines where an unaffiliated, rather than affiliated, provider offers the broadband content service. USF would lower the price of unbundled facilities by the same amount as it would under current rules lower the price of retail services. USF should also lower the cost of broadband common carrier services, such as DSL or FTTx-based data pipes, which should be available on a nondiscriminatory to all broadband content service providers (such as ISPs).

Thus if USF were delivered on such lines, then the ILEC would keep its subsidies and the CLEC or ISP would get access at a reasonable price. There is then no need for CETCs insofar as wireline and broadband services are concerned, USF high-cost support need only be paid once in a given rural locale. (Support for mobility is a separate issue; many current CETCs are of course mobile service providers.)

Spectrum licensing should not create needless scarcity

At 41, the Commission asks,

> We also seek comment on how different regulatory approaches that the Commission has adopted in the past, such as facilitating more efficient spectrum use, developing licensing rules and construction requirements, designating spectrum for licensed versus license-exempt use, secondary markets, cognitive radio, or other polices can ensure efficient and effective access to broadband. For example, what about the adoption of more rigorous buildout obligations for wireless services, such as were recently adopted by the Commission with regard to the 700 MHz band?

Spectrum policy reform is long overdue. While auctions are a more appropriate way to distribute licenses than lotteries (in which the winner typically auctioned off the license privately), the focus on maximizing auction revenue has overwhelmed the Commission's plenary responsibility to manage the spectrum on behalf of the public interest.

Here the major problem is licensing policy. CMRS licenses are generally used for mobile telephone service, which generates the highest revenue for a given unit of network capacity. These licenses are highly concentrated. A few are issued on a CMA basis, which in some rural areas is fairly granular, but much of the spectrum is aggregated into larger blocks which are good for mobility but not well suited to local fixed-wireless operators. And many of these licensees concentrate development in the urban portions of the license areas, leaving fallow spectrum in rural areas.

One solution is to encourage, or require, partial disaggregation of certain licenses. In particular, spectrum licenses that cover both urban and rural areas may be more heavily utilized in the urban areas, where the licensee needs to operate additional frequencies for capacity. In rural areas, the same licensee might only need to use half or less of its allocation, if it chooses to build out there at all. This fallow spectrum could be used to provide rural broadband service.

For example, a 30 MHz (15+15) allocation is sufficient for three CDMA (either CDMA2000 3x or WCDMA) 5 MHz carriers in each direction. In a rural area, it might need to only operate one of these (totaling 10 MHz) in order to handle its requirements. A second 10 MHz might be useful for transition to LTE, but the third is still spare. A rural provider should be able to make

use of the remaining 10 MHz of spectrum to provide broadband in unserved or underserved areas. While leasing of such spectrum is theoretically possible today on a commercial basis, licensees have little incentive to do so.

Another approach is "use it or lose it", which would prohibit the spectrum "banking" that is still common place. Licensees often meet nominal buildout requirements in part of their license footprint but do not serve the whole area, or use their entire licensed spectrum. To some extent this may be done with the goal of reducing competition, behavior which the Commission should discourage, not encourage, even though it may generate higher auction revenues. We note that the 2500-2690 MHz band has very concentrated license holdings, with one carrier often holding licenses to 72 MHz or more in a single market. Worse, these licenses are often put to no use at all. While carriers should be allowed to build footprint, they should not be allowed to block others from competing.

As a general rule, the goal of licensing, including the auction process, should be to make efficient and effective use of the public spectrum. Winning an auction in order to bank spectrum and prevent competition should be prohibited. Auctions should be for the purpose of selecting who *does* make use of spectrum. The Commission should investigate stricter spectrum holding caps in order to discourage this behavior. While the older 45-55 MHz caps are no longer adequate, and new spectrum in the 700 MHz and AWS bands has increased the amount available, today's spectrum holdings are simply too concentrated.

Additional unlicensed or lightly-licensed (nonexclusive) spectrum would also be useful. This has the most potential for low-cost rapid entry by entrepreneurs, and offers equipment manufacturers the most room to innovate. Today's unlicensed allocations are for the most part very busy, but demand for wireless broadband access continues to grow. "White space" is part of the picture, but current rules are very restrictive. If more federal spectrum turned over to civilian use can be made available on an unlicensed or non-exclusive basis, then service is likely to be available to many more users than can be reached via white space.

Define "open" in context of the proper framing

At 47, the Commission asks,

> We seek comment on the value of open networks as an effective and efficient mechanism for ensuring broadband access for all Americans, and specifically on how the term "open" should be defined. For example, should it incorporate access, interconnection, nondiscrimination, or infrastructure sharing principles? The Commission, through its Computer Inquiry proceedings, developed specific nondiscrimination requirements for facilities-based telecommunications carriers,63 although several of these obligations have been scaled back by the courts and by the Commission's revised regulatory framework for wireline broadband Internet access services and other deregulatory measures.64 However, as the regulatory framework for broadband Internet access services changed, the Commission has taken steps to clarify the importance of open networks.65 For instance, the Commission published its Internet Policy Statement establishing four principles "to ensure that broadband networks are widely deployed, open, affordable, and accessible to all consumers."66 More recently, the Commission clarified its authority to enforce those principles and has initiated a proceeding to review broadband industry practices generally.67 In addition, as discussed below, the Commission adopted a requirement for licensees in the 700 MHz Upper C Block to provide an open platform for devices and applications, subject to certain conditions in the 700 MHz auction.68

It must be emphasized that one cannot achieve desired goals of Internet "neutrality" simply by attempting to regulate the behavior of information service providers. As noted above in the framing of this Comment, there must be a clear distinction between carriage and content, wherein the carriage function ("el broadband") must be "open" to a broad range of content providers ("la broadband") who can be regulated by the marketplace. This is *precisely* the concept behind the Commission's Computer Inquiries, which were absolutely essential to the birth of the public Internet. The revocation of the Computer Inquiry rules was a tragic error by the former Commission. It was explicitly *not* mandated by the *Brand X* decision, which merely upheld longstanding Commission precedent that non-common-carriers who provision their own information services do not become common carriers in the process. ILECs are not the same as cable and CAP facility operators and there is no reasonable "level playing field" argument. Certainly the Commission could have adopted policies that retained the protections of the *Computer Inquiries* without violating the broad *Chevron* deference found within *Brand X*.

We see the word "open" as referring to the carriage function, the "basic services" rendered under the guidance of the Computer Inquiries. Through 2005, these could include circuit-based services up through high-speed SONET, packet-based services such as Frame Relay, ATM, and "Ethernet", and consumer access services such as DSL. These basic services – broadband transmission – should be available on an open basis, meaning that at least one common carrier is available to provide the transmission service between customers and a *wide choice* of broadband content services who have *open* entry to the service market via these facilities. This may include a broadband "carrier of last resort" obligation. The obvious holders of this obligation should be ILECs, whose embedded plant gives them the lowest cost, but inasmuch as it may come with access to subsidy money, others may wish to take the role upon themselves.

In the longer term, we recommend that the ideal direction is not to separate carriage and content but to move the split down the stack, and separate facilities from services. As noted above, facilities, not services, should be subsided by the Universal Service Fund. If facilities are structurally (separate ownership) or functionally (separate subsidiary, as in Computer II) separated from the services (including lit transmission services) that ride upon them, then a natural monopoly in facilities (dark fiber and copper) can support a highly-competitive market for transmission (lit fiber and copper) services, which in turn begets a highly-competitive market for higher-layer services. In this case the provider-of-last-resort obligation would most likely flow to the facilities entity.

Some have complained that requiring common carriage would "hamper" private investment. But this need not be the case – the *price* for such service need not be sub-compensatory, it need merely be just and reasonable. Transmission is less competitive than information and content service and thus less subject to pricing pressure. It is the forced vertical integration of the two that hides this fact, and allows transmission owners to sell content at prices impacted by their monopoly power. When transmission is *open* to all comers, its market is widened, the dominant provider's market share grows, and its cost per user almost inevitably falls. This improves its profit margin. The main reason, then, why transmission owners do *not* make wholesale facilities available is that they are hoping to exercise monopoly power over content. This is the reason why the Commission is under such pressure to regulate content.

By the same token we oppose the moves to turn the Commission's "principles" into law or regulations. These have not been properly discussed or investigated, nor scrutinized in the context of a proper rulemaking. To the extent that these are regulation of *content* of Information Service Providers, we believe that they are misguided. The public Internet did not develop *after* a regulatory dictate determined what must be carried, how. It simply evolved in the context of *Computer Inquiries* carriage obligations.

The Internet is not mature. It is not a fixed, stable platform. Its protocol suite (TCP/IP) is already showing major signs of senescence. Some of its backers are pushing a variant protocol, IP version 6, but this direction has met with considerable disagreement and there are numerous reasons to believe that it will not and should not be widely adopted. Yet who will determine the future if content providers such as ISPs are subject to regulation? Does the IETF get rulemaking authority as a QUANGO (quasi non-governmental organization), in violation of the Administrative Procedures Act? Does the FCC become a protocol standards committee? Neither of these is a good solution.

Furthermore, the Internet is subject to constant attack. Viruses, spam, botnets, forged control messages, and malware galore pollute its circuits today. "Neutrality" rules that prohibit ISPs from dealing with these threats will destabilize the net, to the benefit of no one but the criminals.

Separating carriage from content as we have outline above solves the two critical problems that otherwise are mutually exclusive: It allows the Commission and regulators in general to keep their "hands off the Internet" while at the same time ensuring that as much "neutrality" happens as the market can absorb, and that widespread innovation can continue unfettered.

We are attaching an article that describes an opinion of "what it means to be Internet". It offers a definition that provides the flexibility needed to improve and extend service: Internet is a business model, one which can be entirely defined as *a voluntary agreement among network operators to exchange traffic for their mutual benefit*. Regulation of "broadband" content should not require operators to carry traffic on an *involuntary* basis. The PSTN business model is different; it includes the requirement to interconnect and terminate traffic, which necessarily requires the regulation of rates charged for this function (i.e., the terminating monopoly). The two are complementary. While broadband *transmission* belongs as part of the PSTN business model, broadband *content* is not carriage, and regulating it as such is a slippery slope.

Privacy must be respected

At 59, the Commission asks,

> …What are consumer expectations of privacy when using broadband services or technology and what impact do privacy concerns have on broadband adoption and use? We also note that certain broadband providers have purchased the behavioral advertising86 services of companies that advertise an ability to "deliver[] the most actionable consumer intelligence by extending [those companies'] reach dynamically to encompass the ever-growing network of sites that consumers visit."87 These companies track the webpages customers visit, the searches they perform, and the ads they click, among other information. Consumers may also be aware of the technological ability that broadband providers have to perform functions such as deep packet inspection.89 What is the impact of this type of activity on consumers' willingness to use broadband services? We seek comment on how the Commission should treat issues such as deep packet inspection and behavioral advertising in developing a national broadband plan and whether there are issues related to other types of information connected with the provision of broadband services that the Commission should consider.

To the extent that any transmission or service is common carriage, wiretapping laws should provide the basic guidance for how content is tracked. These generally provide for a high expectation of privacy, with carriers allowed to perform service observation for maintenance purposes but not divulge or otherwise make commercial use of information obtained in doing so. Even private carriage should generally be subject to similar rules – when a customer purchases transmission service, the content of the transmission remains the customer's business. This should be the rule for broadband transmission, just as it is with the PSTN. In the case of a packet-switched protocol, header information (technically, protocol control information; it could, for instance, include trailers) is potentially used by the network in order to transmit the message, but the defined payload should remain sacrosanct. In general we note that today's widespread TCP/IP protocol was not designed as a transmission service protocol but to be a higher-layer service protocol; transmission protocols should be ultimately left to the market but in today's context typically include TDM circuits, Frame Relay, Asynchronous Transfer Mode, and

"Ethernet" switching. However we do not exclude the fact that non-Internet transmission services also use TCP/IP, as it has become rather ubiquitous.

In the case of the broadband services (ISPs, etc.) carried atop the transmission function, the situation is somewhat different. It is our expectation that if there were truly an open, fully competitive market for these services, rather than a duopoly linked to physical carriage, then the types of abuse that have occasionally been seen or alleged would not thrive. Customers would simply go elsewhere. However, this does not excuse violations of privacy. At the same time, if a service is "information" and not carriage, then almost by definition one would expect that the service provider is performing more than bit-level or packet-level relaying. Even editorial judgment is acceptable for some kinds of service.

Hence the two most important rules for broadband service providers should be respect for customer-private information and full disclosure of policies. The content of ISP transmissions, even on an aggregated or anonymized basis, should not be divulged to third parties, even affiliated arms of the ISP, unless the customer has explicitly opted in to such a plan, wherein the amount of information to be divulged and to whom is clearly explained in advance and cannot be expanded without a further opt-in.

The term "deep packet inspection" (DPI) has been applied rather broadly, and it is useful to distinguish between different forms Some people believe that the TCP/UDP header is part of the payload and thus any application that looks at it, including Network Address Translation, is DPI. This is a stretch: TCP and IP were once one protocol, and when the two were split in the 1970s, the closest thing it has to a multiplexing function (port number) was erroneously put in TCP, when it more appropriately an IP function. Hence NAT performs a *syntactic* change in TCP (port number swapping) while maintaining the basic *semantics* (what service is being reached) of an address-port pair. This is hardly DPI. Similarly, some application protocols are defectively designed in that they put the numeric IP address inside the application-layer header. These too need to be adjusted by the NAT, an action that does not really invade anyone's privacy.

Moving slightly deeper, there are applications (such as BitTorrent) which create parallel TCP sessions, which violates the basic "fairness" mechanism of TCP that treats each TCP flow equally. This can be detected by DPI and in turn that can be used to throttle Torrent traffic in

times of network congestion. This again does not require the divulgence of customer information and does not *prevent* any usage of any protocols, and thus can be seen as acceptable practice if divulged. (Of course Torrent users believe that this is "discriminatory", but that call is best left to the market. Users *other than* Torrent users receive better service when Torrent's approach to dealing with TCP capacity rationing is thwarted.)

Likewise, if an ISP chooses to offer a restricted service via its Terms of Service, and makes clear that (for instance) it does *not* permit file transfer servers (these include so-called "peer to peer" applications, which include file transfer clients and servers) to be operated at subscriber locations, then "policing" mechanisms that are used to locate rule violators are not a violation of privacy or customer expectations. They are *not* "neutral" but there is no need for such neutrality, as lower-layer broadband transmission *should* be available on a neutral basis, allowing many different ISPs to market plans with different rate structures and terms of service, optimized for different types of users.

On the other hand the current *absence* of multi-provider competition for ISP services *invites* transmission owners to abuse DPI for commercial purposes. DPI vendors routinely advertise the ability of their systems to "monetize" Internet traffic. This could include charging for email by the message, charging different prices to reach different web sites, blocking or charging usage fees for VoIP, blocking or charging usage fees for different streaming media, and even "taking a cut" of e-commerce transactions. Such practices already exist, more or less, among CMRS wireless carriers, which is one reason why they should not be viewed as full-scale broadband competitors. But these practices would probably be starved out of existence by fair competition at the service layer. Until such competition is restored (i.e., by repairing the separation of content and carriage), then it may be necessary to dictate that ISPs who have significant market power and own their own transmission facilities *may not* engage in such anti-consumer practices. Otherwise, "broadband" will degenerate into "Fat Wasteband, Broadband Internet's Evil Twin". That "Minitel on steroids" approach is not a suitable one for the nation's future.

Attachment:

What does it mean to be Internet?

Fred Goldstein, Ionary Consulting, June 2009

It's a simple question and it can be read two ways. One way is to ask just what constitutes the Internet, or makes something a part of the Internet. The other way is to ask whether and why it matters. Those are important policy questions in today's world, where "Internet" is still a magic word, beloved of many policy makers but still widely feared.

The first question is critical. Defining just what the Internet is not simple. Justice Potter Stewart once explained that he couldn't easily define obscenity, "but I know it when I see it". We think we know what the Internet is, but to *define* something is to draw a limit as to what the word does and doesn't refer to. Examples don't count; they are just illustrations. We know that the Internet today contains millions of computers in a worldwide network-of-networks that uses the TCP/IP protocol stack to provide a range of applications, but the Internet isn't defined by its protocols or applications any more than "horse" is defined by the Kentucky Derby or Alpo.

First, a definition

In fact, the Internet we are using today is not the final product. It's a prototype, a lab experiment run amok. It's a complicated experiment, one that goes beyond computers and cables and fibers and into whole methods of organizational development, standardization and financing. But it's not the final product. That only emerges after answering the question with two meanings, what does it mean to be Internet? A real definition would be useful, and exactly how that definition is crafted could have many implications. So here's a definition that is no more or less encompassing than necessary:

Internet (n.) A voluntary agreement among network operators to exchange traffic for their mutual benefit.

This definition is intentionally terse, a distillation of the concept to its essence. It simply defines a business model for interconnecting networks. It does not specify a protocol. TCP/IP has been in general use since the early 1980s, but it's running out of steam, and a replacement protocol or two may come along soon. Calling something "IP" shouldn't convey an exclusive right to be

considered as Internet. By the same token, use of an Internet Protocol such as IP does not make something Internet. IP is already used within the telephone network and other non-Internet applications. Certainly the telephone industry's IMS (IP Multimedia Subsystem) is anything *but* Internet! IP is a tool. It's a tool designed *for* the Internet, but not all usage of IP is Internet.

The definition does not include any reference to the current global big-I "Internet". That is an Internet, but it need not be the only one. Today's prototype demonstrates the Internet business model, but even now, some of its participants are not really happy about it.

The definition does not include a reference to the Internet Engineering Task Force, the Internet Society, ICANN, the United States government, or any other organization. They are all participants, and they all helped create it, but their continued participation doesn't define the model. It's not an exclusive club. Anyone can play.

What the definition does specify: It is *voluntary*. This is a very important part of the business model, one that is unexpected in the telecom industry, which is characterized by historical monopolies, but common in other forms of commerce. Indeed, Internet is basically what you get when ordinary commerce takes place in the form of transmitted information, rather than hard goods.

Network operators don't have to join in, or play by anyone's regulations. An Internet is basically a set of contractual, usually bilateral, agreements. So if a network operator doesn't want to allow another one to connect, it doesn't *have* to. It doesn't *have to* accept any traffic it doesn't want, whether it deems it spam, malware, or simply excessive. It doesn't have to pay anyone their asking price; it can just walk away – there is no mandatory interconnection between ISPs. There is no tariff with fixed prices open to all comers. And thus there is no bright-line distinction between network operator and customer, no formal definition of who is or isn't an ISP. Peering – the interchange of traffic between ISPs at no cost, the privilege that distinguishes retail customers from ISPs – is not mandatory, but reflects the *mutual benefit* of the *exchanage of traffic*. And if that mutual benefit involves a little cash in either direction, fine. That's what voluntary is about.

The telecom industry follows a contrasting business model

Contrast this to the traditional telecom industry, which operates the Public Switched Telephone Network (PSTN). That's based on a regulated public utility business model, one that really goes

beyond plain old telephone service (POTS). The PSTN has licensed service providers who are *common carriers*. They file tariffs which set uniform rates for all takers. Most PSTN providers are also all too happy to bill for every little transaction, an "annoy the customer" feature that festers in a monopoly environment.

PSTN Protocols are often specified by authorities, or by large Incumbent carriers who have the type of monopoly power that doesn't occur in the Internet. There's a bright line distinction between a PSTN customer ("subscriber") circuit and an intercarrier connection. Interconnection between carriers is mandatory – they *must* accept each other's traffic, but the rate they are allowed to charge each other is also set by regulatory authorities, via tariff or regulated contract.

This isn't really a bad business model, since the PSTN helps provide "universal service" in the sense of one worldwide network that almost anyone can belong to at a predictable price. It is well suited to dealing with monopolies and with activities where competitive forces are weak. But it's not very flexible, and it is prone to both monopoly abuse and regulatory friction.

The distinction between Internet and PSTN thus does not hinge on whether the traffic in question is voice, data, video, or anything else. It doesn't hinge on whether the switching is based on packets, circuits, or for that matter smoke signals. Those are just the tools. It hinges on the business model. The Internet business model, *when* it works, can be applied outside of the "big-I Internet". And the PSTN can make use of the same tools, such as packet-switching protocols, without becoming Internet. **The two are largely complementary.** To some extent they compete, but they have different strengths and weaknesses, and both sides need to learn to work better with the other. The highly-competitive Internet model only thrives when given access to regulated PSTN carriage, which needs to treat it as bit-neutral customer payload.

Implications of the voluntary model

So what else does it mean to be Internet, in that second sense, now that we've defined it? For one thing, it means that orthodox "network neutrality" is literally impossible to apply to ISPs, as that means *involuntary* interconnection. At that point the ISPs have been swallowed into the PSTN!

The problem is that the Internet sometimes works so well, albeit by brute force, that people have started to use it as a substitute for the PSTN and even for cable television, yet another business

model with its own regulatory baggage. Its voluntary, free-market business model allows it to be cheaper than the PSTN for some applications, since the PSTN lacks sufficient competition, and still heavily taxed and cross-subsidized in order to fund "universal service" and to achieve other political goals. And no good deed goes unpunished. Even worse, some are calling for regulation of Internet content providers (web site operators), to regulate what they charge or to whom they must make their proprietary content available.

Not only are these ideas tantamount to regulating the press, but they're like telling publishers what they can charge and to whom they must sell their wares. Internet is about content, not carriage. That's why it is classified in the United States as an "information service". A free press does not mean one in which all publications must be given away free-as-in-beer. Democracy is the right to own a press, not to dictate what others publish. And the lack of access from the subscriber *to* an unrestricted choice of ISPs – the current cable/telco (think of *Pravda* and *Izvestia*, the two national newspapers of the old Soviet Union*)* duopoly – is the real threat to network democracy.

The PSTN, in contrast, *must* be content-neutral; because that's what its role, common carriage, is all about. We almost had a PSTN able to provide high-speed packet-switched data, video, and other services – that's why ATM (Asynchronous Transfer Mode, a networking protocol still widely used within DSL access networks) was invented over two decades ago. But it failed because the billing-happy PSTN business model couldn't accommodate it, and the Internet stepped in to fill the void.

The post-2005 US regulatory model, which treats the Internet as a vertically-integrated stack right down to the wire or radio waves that carry it, is literally self-contradictory. Something can't be Internet and PSTN at the same time. It doesn't work. The resulting regulatory friction is already proving to be incendiary. Yes, ISPs can be self-provisioned non-common-carriers, but the telecom industry can't be treated as Internet just because it also provides retail ISP services. Indeed the public Internet only became possible because of regulations the FCC had imposed well before 1984's AT&T divestiture. One was a regulatory requirement that carriers allow sharing and resale of its lines. The other was a requirement that content-based ("enhanced") services be treated as neutral payload, with PSTN carriers having to treat their own enhanced services the same as their competitors'.

Nobody's in charge here, move right along

Another thing that being Internet means is that *there is no authority*. PSTN operators are beholden to government regulators; the Internet isn't. ICANN, the Internet Corporation for Names and Numbers, is merely a voluntary organization. It coordinates names and numbers on the "big-I Internet", but ISPs are free to ignore them and pick IP addresses of their own choosing. Of course since interconnection is voluntary, other ISPs are equally free to refuse interconnection, or they can choose to create their own workarounds such as encapsulation or translation. This sort of thing is unlikely to happen much, but it provides a check on organizational behavior. ICANN is, in effect, elected by ISPs who choose to use it.

This lack of authority extends even more strongly to the IETF. This voluntary organization dates back to the days when the Internet was the private domain of the government and its selected partners. It uses a "consensus" technique to create Internet standards. Consensus generally boils down to how many people show up in the room. Organizations and countries don't vote; even individuals don't vote. The loudest position wins. This tends to favor large organizations who have a vested interest and enough geeks and money to make sure that they're well-represented at meetings. It has not proven, however, to be effective in maintaining consistently high quality.

The IETF, and its overseer the Internet Architecture Board, lost sight of the ball about 15 years ago when they adopted IP Version 6 as the "next generation" Internet protocol. IPv6 fails almost every test for a proper next generation protocol. It does not really address multihoming, multicast, or mobility, three weaknesses of IP that have been understood for decades. It doesn't handle inter-provider scaling and it lacks backwards compatibility. It is merely a costly, crude fix for a perceived shortage of IP addresses. But the IAB and IETF are unwilling to admit to having made a mistake, so their current answer is to shout louder, even as most ISPs and Internet users correctly ignore them.

So a new Internet may well evolve entirely outside of the purview of the IETF and IAB. It could take shape among service providers, equipment vendors and users. It could take place without even being noticed by the government or big operators. Or the *voluntary* standards bodies could catch up. Humans are social animals with an urge to communicate. And thus Internets will continue to be built.